CONGRESSIONAL BUDGET OFFICE
U.S. Congress
Washington, DC 20515

Douglas Holtz-Eakin, Director

February 28, 2005

Honorable Joseph I. Lieberman
Ranking Member
Subcommittee on Airland
Committee on Armed Services
United States Senate
Washington, DC 20510-6050

Dear Senator:

In response to your request, the Congressional Budget Office (CBO) has reviewed several issues that have arisen regarding choices that the Department of Defense (DoD) has made in implementing the Global Information Grid Bandwidth Expansion, or GIG-BE. CBO's review, which is provided in the attachment to this letter, focuses on the following issues:

- Whether the GIG-BE will provide substantial additional capacity to transport classified traffic relative to DoD's existing communications networks;
- Whether the GIG-BE will as secure as DoD's existing networks; and,
- Whether the GIG-BE will provide the quality of service that users of DoD's existing networks have come to expect.

CBO's review focuses on the GIG-BE, not the overall Global Information Grid (of which the GIG-BE is a key component). According to DoD's current plans, the overall Global Information Grid will encompass numerous programs involving sensors that collect information, computer systems that process information, and systems that transmit information to and among combat units on the move, as well as to and among fixed facilities (the latter being, at least initially, the primary purpose of the GIG-BE).

If you would like further details, I would be pleased to provide them. CBO's staff point of contact is J. Michael Gilmore, who can be reached at (202) 226-2917.

Sincerely,

Douglas Holtz-Eakin

Attachment

Issues Associated with the Global Information Grid Bandwidth Expansion

February 28, 2005

Congress of the United States
Congressional Budget Office

Contents

Tables

Figure

Introduction and Summary

The Department of Defense (DoD) is undertaking a major initiative to create a communications network called the Global Information Grid (GIG). The GIG is intended to be a single information-sharing network with multiple levels of security for use by DoD and possibly other elements of the national security community, including the intelligence agencies. It is intended to satisfy all operational and business needs for transmitting and sharing information among elements of DoD in both peacetime and wartime. According to DoD, the GIG will be key to implementing the department's plans for what it calls network-centric warfare, which focuses on rapidly sharing information among all levels of military command to win battles. Although the content of the GIG continues to evolve, it currently encompasses numerous programs that will be executed over many years to address needs for both tactical (associated with combat units on the move) and strategic communications. Portions of the network have already been implemented and are being used on a daily basis by elements of DoD.

This Congressional Budget Office (CBO) document focuses on one part of the GIG: the Global Information Grid Bandwidth Expansion (GIG-BE). That project involves upgrading the transmission capacity of some of the busiest parts of DoD's existing communications networks. Final testing of the GIG-BE is scheduled to occur in 2005. However, the technological approaches that DoD has taken to carry out the expansion have raised questions. In particular, some observers have expressed concerns about the degree to which the GIG-BE's capacity, security, and quality of service will represent improvements over DoD's existing networks. This CBO review addresses those concerns by examining some of the choices that DoD has made for the GIG-BE and their implications.

Some Key Concepts About Communications Networks

Communications networks are composed of many elements of hardware and software that provide services to users at various levels.[1] The GIG-BE is intended to provide high-capacity communications linking DoD users at locations worldwide.

Understanding DoD's choices for implementing the GIG-BE requires some familiarity with concepts such as the Internet protocol (IP), a method for transmitting many different types of data, segmented into packets, between computers connected to a network; the fiber-optic cables, or other links, that compose a network's physical connections; and other devices, called routers or switches, that send messages and data along the links. At the edges of any communications network (the places where data move into or out of the network) are local area networks (LANs) or wireless (radio) networks, to which user devices, such as telephones and computers, are attached. This review does not focus on devices at

1. Those elements are discussed in more detail in Appendix A.

the edges of the GIG-BE but rather on the backbone, or core, network, which is composed of the links, routers, or switches that move data to and from the many user devices connected to the GIG-BE at its edges.

Communications networks are commonly described in terms of a model—the Open Systems Interconnection Reference (OSIR) Model—that was developed in 1983 by the International Organization for Standardization. That model offers a structured approach for defining and developing the functions and protocols that enable computer systems to connect to one another. The OSIR model is composed of seven layers, with each successive layer providing a service to the layer above it. Data originate at layer 7 (the application layer) on a device such as a personal computer and then flow down to layer 1 (the physical layer), where the data are transmitted as a stream of bits across a physical connection, such as fiber-optic cables, between elements of the network. The data then flow back up to layer 7 of the systems at their destination. (For descriptions of all of the layers in the OSIR model and their functions, see Appendix A.)[2]

This review focuses on the protocols and hardware being used in two layers of the GIG-BE: layer 2 (the data link layer), which manages transmission for the network layer and controls the flow and delivery of data; and layer 3 (the network layer), which handles routing and congestion and establishes and tears down connections between devices that communicate over the network.

Issues Associated with DoD's Decisions for the GIG-BE

The GIG-BE is similar in several aspects to communications-network projects that have been undertaken in the civilian sector, and the DoD has used those similarities as part of its rationale for the technology chosen to implement the GIG-BE. Specifically, in carrying out the expansion, DoD has opted to:

- Use dense wave division multiplexing, a commercially available optical transmission technology for layer 1 that provides much greater capacity than that of DoD's legacy (preexisting) networks;[3]

- Employ commercially available technology at layers 2 and 3 that many observers say is well suited for providing converged services using the

2. Not all networks use all seven levels of the OSIR model, which is both general and idealized. In particular, the Internet protocol architecture uses three layers: application, transport, and network.

3. Unless otherwise noted, the metric used for a network's capacity is its bandwidth. Bandwidth is usually measured in terms of the number of bits per second that can be transmitted, typically in gigabits (billions of bits) or megabits (millions of bits) per second. A bit is a piece of binary data; that is, it is either zero or one.

Internet protocol.[4] In particular, the GIG-BE will use multiprotocol label switching (MPLS) to forward packets of data through its core.[5] MPLS's forerunners have been in existence since 1995, and the Internet Engineering Task Force published guidance for the use of that technology in January 2001.

- Perform encryption and create virtual private networks (using MPLS) at layer 3 to separate classified and unclassified traffic—rather than at layer 2, as in most of DoD's legacy networks.[6]

- Support the transport of traffic sent over DoD's legacy networks. As noted above, legacy traffic that is classified is generally encrypted using devices that operate at layer 2 (and sometimes layer 1). For transmission within the continental United States, the GIG-BE will provide additional capacity for classified traffic using encryption performed at layer 3—called high-assurance Internet protocol encryption—which is being developed by the government and is based on a commercial analog, called Internet protocol security.[7]

The issues that have been raised about the GIG-BE's implementation deal with the last three choices listed above. In particular, critics have maintained that because the GIG-BE will use high-assurance Internet protocol encryption, it will provide only limited additional capacity to transport classified traffic compared with the capacity of DoD's legacy networks. Moreover, some critics charge that the GIG-BE will not be as secure as those legacy networks and will not provide the quality of service that users of those networks have come to expect.

---

4. "Converged services" refers to networks that simultaneously carry voice, video, and data traffic. DoD regards the ability to support converged services using the Internet protocol as key to realizing its goals for network-centric warfare.

5. MPLS is not formally part of the Internet protocol. It is typically used to transport traffic over a core network linking groups of subscribers; it is not used directly by those subscribers. Some network engineers see MPLS as operating between layers 2 and 3 because it can be used in conjunction with other layer-2 technologies.

6. Encryption is the process of encoding or decoding information to protect it from being read by people not authorized to see it. A key is required to encode (encrypt) and decode (decrypt) the information. A virtual private network is composed of a set of users at different sites, often geographically dispersed, that communicate with each other using a common core network, such as the GIG-BE.

7. Internet Engineering Task Force, *Request for Comments: 2401—Security Architecture for the Internet Protocol* (November 1998), available at www.ietf.org/rfc.html. For transmissions outside the continental United States, the GIG-BE will use encryption performed at layer 1, also called bulk, or link, encryption.

CBO's Conclusions

This review examines each of those claims, using information provided by the Defense Information Systems Agency about the design and capabilities of the GIG-BE; standards and other informational material published by the Internet Engineering Task Force, the ATM (asynchronous transfer mode) Forum, and the MPLS and Frame Relay Alliance; and other source material, including assessments performed by the National Security Agency.[8] On the basis of those sources, CBO's review concludes the following:

- Initially, the GIG-BE's capacity for transporting encrypted traffic will be about double that of an analogous legacy network. Under current plans, that capacity will grow, and the GIG-BE will eventually provide about 10 times the encrypted capacity of a similar legacy network.

- Debate continues in the national security community about the benefits and drawbacks of the technological approaches used to ensure the confidentiality of classified information in the GIG-BE versus the approaches used in DoD's legacy networks. However, an assessment by the staff of the National Security Agency concludes that there is no reason to prefer existing approaches to those used in the GIG-BE.

- The GIG-BE will not initially provide all of the quality-of-service features of DoD's legacy networks, but it has the potential to do so eventually.

The Global Information Grid and Its Bandwidth Expansion Project

The Global Information Grid is a communications network being developed by the Department of Defense that will be capable of processing and transmitting information at multiple levels of security classification. The GIG will provide a single network that will enable users throughout DoD—and potentially the intelligence agencies—to share information.

DoD uses and maintains a host of separate computer networks. Many are local area networks that provide services over limited geographic areas (such as inside

8. The ATM Forum and the MPLS and Frame Relay Alliance are international nonprofit organizations formed to promote the use of their respective technologies through a variety of means, including development of interoperability specifications. Membership in the organizations consists primarily of companies that sell equipment and services utilizing the technologies that the organizations promote. The Internet Engineering Task Force describes itself as "a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. It is open to any interested individual."

military bases and other defense facilities). Generally, those LANs need to be linked so that users located at facilities across the United States and at other sites worldwide can share data. To accomplish that linkage for DoD users, the Defense Information Systems Agency operates and maintains the Defense Information System Network (DISN).[9] The DISN consists of a number of elements. For example, it includes the DISN Asynchronous Transfer Mode System (DATMS), which transports both classified and unclassified Internet protocol traffic using the layer-2 technology that is part of its name—asynchronous transfer mode.[10] Classified Internet traffic is transported by the Secret Internet Protocol Network, which uses the DATMS.

The Defense Information System Network is composed of about 700 nodes worldwide and the telecommunications links used to transmit data among those nodes.[11] Some of those links, known as private lines, are entirely owned and maintained by DoD. Others, called public leased lines, are leased from commercial telecommunications providers.[12]

The DISN's nodes fall into three types. Class A nodes have a capacity of up to 10 gigabits per second (Gbps); class B nodes have a capacity of up to 2.5 Gbps; and class C nodes handle 622 megabits per second (Mbps) or less. The DISN now includes 16 class A nodes and about 80 class B nodes, with the remainder being class C or smaller.[13]

The Global Information Grid Bandwidth Expansion project is upgrading the transmission capacity of approximately 100 nodes of the DISN around the world (and

---

9. The Defense Information Systems Agency also provides telephone services though the Defense Switched Network, which is part of the DISN.

10. Some network engineers now view asynchronous transfer mode as operating at "lower-level 3" rather than at layer 2.

11. A node is a geographic site where the switches, routers, and other electronic devices that receive and transmit information over communications links (often fiber-optic cables) are located.

12. Most of the GIG-BE's links are public leased lines, although some new links, including ones that address new network security requirements, are being purchased by the GIG-BE. Therefore, the GIG-BE is also laying some private lines. Because of the large excess capacity that exists in commercially available fiber-optic cable around the globe, DoD can acquire public leased lines—particularly ones that connect major population centers—relatively cheaply. Anticipating much greater demand for bandwidth than occurred, telecommunications companies laid a huge amount of fiber-optic cable worldwide in the late 1990s. Estimates of the amount of that cable currently in use range from 2 percent to 5 percent. The unused capacity is known as dark fiber. Much of the new capacity provided by the GIG-BE "lights up" commercial dark fiber.

13. The intelligence community operates and maintains another network similar to the DISN that has some nodes in common with the DISN.

Table 1.

The Procurement Budget for the GIG-BE

(In millions of dollars)

	2002	2003	2004	2005	2006	Total
Fiber Upgrade[a]	0	213	185	0	0	398
Switches, Routers, and Other Devices	0	231	159	0	0	390
Integration[b]	0	48	19	10	0	77
Total	0	492	363	10	0	865

Source: GIG-BE procurement exhibits supporting the President's budgets for fiscal years 2005 and 2006.

a. A consolidation of line leases, link upgrades, and facilities upgrades.

b. Includes transition and management costs.

associated links) that handle the largest amounts of traffic.[14] Those upgrades are expected to be completed by the end of 2005. Integration of the GIG-BE with the remaining nodes of the DISN will take place later under the new DISN Access Transport Services program.

Costs of the GIG-BE

DoD's budget for buying equipment and links for the GIG-BE has totaled $865 million through 2005 (see Table 1).[15] In assembling the links that compose its core network, the GIG-BE program has taken advantage of the large amounts of dark fiber (unused commercial fiber-optic capacity) that can be leased quickly and cheaply. The program has also made extensive use of commercially available hardware, such as routers and optical transmission equipment. The high-assurance Internet protocol encryptor devices being used in the GIG-BE are not commercial equipment; they are being developed as part of government-sponsored programs other than the GIG-BE.

14. Of the 100 nodes in GIG-BE, two are nodes being created by the GIG-BE program that are not already DISN nodes. In the continental United States, the GIG-BE will consist of 66 service-delivery nodes.

15. A number of other budget accounts controlled by the Defense Information Systems Agency contain funding for activities supporting integration and operations of the GIG-BE, as well as similar functions for other defense communications and computer networks.

Asynchronous Transfer Mode and Multiprotocol Label Switching

Like a number of existing commercial networks, the DATMS and several other networks currently used by the national security community to transport data were built during the 1990s and employ asynchronous transfer mode (ATM). Some of those networks have been upgraded in the past year using high-capacity combination switch-routers capable of supporting both ATM and other methods for forwarding traffic, including multiprotocol label switching.[16] The GIG-BE is using MPLS rather than ATM to forward data. Both methods have common benefits, but MPLS can also offer some advantages in efficiency and simplicity. As a result, MPLS appears to be gaining in popularity in the commercial market.

Benefits and Drawbacks of Connection-Oriented and Connectionless Services

Under the Open Systems Interconnection Reference Model, services provided to the network layer (layer 3) for transporting data are characterized as being either connection-oriented or connectionless. Connection-oriented services establish a path (or connection) from source to destination before transferring data; in addition, they generally reserve paths and network capacity for forwarding certain kinds of traffic. Connectionless services, such as the Internet protocol, transfer data from source to destination without first establishing a predetermined path and generally without reserving capacity for carrying specific kinds of traffic.

Both ATM and MPLS are connection-oriented services that are employed in conjunction with networks using the connectionless Internet protocol. When used to forward Internet traffic through a core network, ATM or MPLS reads the Internet header on a packet of data at the core's edge to determine the path through the core that the packet should traverse. A short header is appended to the packet specifying that path, and only that short header, rather than the packet's longer Internet header, is processed by the ATM switches or MPLS routers in the network's core. Once ATM or MPLS establishes a path, data flowing over that path generally arrive in the order in which they were sent.[17]

Such connection-oriented services offer an advantage over connectionless services in reducing congestion in a network. It is difficult for connectionless services to route traffic away from congestion points, and in some instances, connectionless services can cause congestion to occur despite the existence of unused or underutilized capacity in the network. In contrast, predetermined routes designed to

16. "Marconi rings up $33 million sale to U.S.," *Pittsburgh Business Times*, June 8, 2004, available at www.bizjournals.com/pittsburgh/stories/2004/06/07/daily14.html.

17. On occasion, however, both ATM and MPLS transmit data out of order.

mitigate congestion can be built for traffic using a connection-oriented service, as can backup routes onto which traffic can be switched automatically when primary routes fail.

Connection-oriented services can be less efficient than connectionless services, however, because they reserve capacity that can sometimes be underutilized. Moreover, both types of services can fail to deliver data—in the case of connection-oriented services, because not enough capacity was reserved to carry the traffic that materialized, or, in the case of connectionless services, because of transient congestion.

Technical Issues Involving ATM

In the early 1990s, Internet service providers had core networks composed of Internet protocol routers connected with lines that supported 45 Mbps of capacity—a "routed core." Those networks had drawbacks, however: the routers had limited capacity and were relatively slow, and the connectionless nature of Internet protocol routing made traffic engineering (routing traffic so as to optimize the use of the network's resources) difficult. As Internet traffic increased, so did the need for traffic engineering.

Beginning in the mid-1990s, the volume of Internet traffic grew to the point where Internet service providers sought solutions other than a routed core. One solution that became available was to build the core using faster (relative to the then-available routers) ATM switches supporting speeds of 155 Mbps to 622 Mbps and to "overlay" the Internet protocol edge network on that ATM core.[18] Those faster ATM switches could make better use of the higher-capacity fiber-optic links that were becoming available and that Internet service providers were using to meet increased demand. ATM switches in the network's core were used to establish "virtual circuits" to provide the necessary connections among the Internet protocol routers at the network's edge. The paths used by the virtual circuits to traverse the network's core, their capacity, and their other attributes could be tailored to give users a predetermined level of service and to make the best use of the network's resources. Backup circuits could also be configured for use in the event that a failure affected the network's primary circuits. The use of fast ATM connection-oriented service in the core gave Internet service providers the ability to manage their networks and provide guaranteed service that the connectionless routed core did not. DoD used ATM switches to build its DATMS network around 1998.

18. ATM is called asynchronous because it is structured to handle traffic sent by users at nonconstant rates. Synchronous transfer provides constantly repeating time slots for users to send traffic. If no traffic is sent, the time slot is not available for another user's traffic. Thus, for traffic that is not sent at a constant rate, asynchronous transfer can be more efficient than synchronous transfer.

Configuring ATM virtual circuits can be a complex task, however. It requires special skills and is often performed off-line (not as part of a network's ongoing operations) using special-purpose configuration utilities (or software packages). One way in which ATM has been used in conjunction with the Internet protocol—and in which it is used in the DATMS—requires that a full mesh of virtual circuits be established that permanently connect all of the Internet routers at the network's edge. If there are *N* such routers, a full mesh of virtual circuits connecting them consists of $C(N,2) = N*(N-1)/2$ linkages.[19] That requirement—referred to as an N^2 scaling problem—has been cited as a drawback to using asynchronous transfer mode.[20] However, it is not always necessary to establish a full mesh of ATM virtual circuits to form a backbone connecting Internet protocol routers. In addition, not all ATM networks are configured using permanent virtual circuits that must be manually established or changed as the network's configuration changes. Some ATM networks use virtual circuits that are established and de-established dynamically and automatically.

Another set of issues arises because ATM segments and then reassembles the variable-length Internet packets generated by the overlying Internet protocol network into a series of cells, each 53 bytes long (48 bytes of data and a five-byte cell header).[21] On the one hand, that practice is part of the reason that ATM switches operated faster than mid-1990s routers—fixed-length cells were then easier to process than variable-length packets. It is also the reason that ATM is particularly good for transporting voice traffic on lower-speed links—smaller-size packets carrying voice will not be delayed because they are awaiting the processing of much longer packets. On the other hand, the operating speed of the hardware used to segment variable-length data packets into fixed-size ATM cells for transport and to reassemble those packets at their destination is a potential bottleneck. Interfaces that convert data for optical transmission are currently available at speeds of up to 10 Gbps. But the hardware commonly used today to segment and reassemble ATM cells operates at speeds of 622 Mbps to 2.5 Gbps. (Hardware operating at 10 Gbps is available but is not widely deployed.)[22]

---

19. The notation *C(N,M)* refers to the combination of *N* things taken *M* at a time.

20. Scaling refers to how the complexity of a network increases as the number of nodes, routers, switches, and other hardware devices in the network increases.

21. A byte is eight bits. The 48-byte cell size for data was adopted by the International Telecommunications Union's Telecommunications Standardization Sector as a compromise between a longer size that would have been more efficient for transporting traffic that is not sensitive to delay (such as transferring large data files) and a smaller size that would have been superior for transporting traffic (such as voice) that is sensitive to delay.

22. In 2000, some observers believed that commercial deployment of segmentation and reassembly hardware operating at 2.5 Gbps would be a challenge and that 10 Gbps equipment might never be feasible commercially.

Efficiency is another issue associated with the use of ATM, related to its use of fixed-length cells. Efficiency is defined as the ratio of bytes of users' data transported to total bytes transported, including overhead associated with network operations. (All networks, no matter how configured, impose overhead.) Because each ATM cell includes a five-byte header—specifying, among other things, the virtual circuit that should be used to transport the cell—the efficiency of ATM can be no greater than 48/53, or 91 percent. But because many of the Internet protocol packets generated are small, and most are not of a length that is evenly divisible by 48, using ATM imposes additional overhead. For example, a 50-byte Internet protocol packet would require the use of two ATM cells for transport, implying an efficiency of (at most) 50/[2*(48+5)] = 47 percent. When averaged over a typical distribution of Internet protocol packets, ATM efficiency is about 80 percent, meaning that 20 percent of the bandwidth in an ATM core is not available to transport users' data. (Network overhead and efficiency are discussed in more detail below in the section on capacity.)

Technical Issues Involving MPLS

Multiprotocol label switching was first developed in the mid-1990s as a means of building core networks that incorporate the advantages of connection-oriented services such as ATM without the need for ATM's unique set of hardware and operational skills.[23] MPLS uses label switched paths—analogous to an ATM virtual circuit—to transport traffic through a network's core. But unlike ATM, MPLS uses extensions of layer 3 Internet protocols to establish its label switched paths. Internet protocol routers that are MPLS-capable (called label edge routers) operate at the edge of a network and append an MPLS label to an outgoing variable-length Internet protocol packet. That label, which is four bytes in length, designates, among other things, the label switched path that the packet should traverse though the core. Routers in the network core (called label switching routers) determine how a packet should be forwarded to its destination solely on the basis of its MPLS label; the routers do not examine a packet's longer (at least 20 bytes) Internet protocol header.[24] That process is analogous to the way in which ATM switches use the ATM cell header to forward cells through an ATM core.

---

23. Some observers, including marketing engineers for MPLS equipment vendors, argue that ATM has disadvantages in terms of complexity and cost because it requires having personnel who are skilled in the operation of both Internet-protocol-based edge networks and ATM-based core networks. Some engineers familiar with the design and operation of ATM-based core networks discount that argument, however.

24. MPLS allows for an arbitrary number of four-byte labels to be appended to a packet. Typical applications use several labels. The ability to use many labels can reduce the number of traffic flows that the core label switching routers have to process, which improves the scaling properties of the core network.

Originally, the use of a shorter header, which could be processed more quickly, was motivated by slow router speeds. Now, however, high-speed routers that exceed the performance of some ATM switches are available, which eliminates that concern. Nonetheless, using MPLS in the core provides the traffic-engineering features of a connection-oriented service without the need to buy and operate the separate infrastructure used by ATM. Also, because MPLS does not segment packets into fixed-length cells, it can be about 20 percent more efficient than ATM in situations that do not require traffic to be encrypted. (The section below on encryption provides more detail about the efficiency of ATM and MPLS core networks for transporting encrypted and unencrypted traffic.) Moreover, the MPLS core does not exhibit the N^2 scaling of a full-mesh ATM core, because the number of label switched paths that core routers must contend with is proportional to the number of label edge routers; that is, their number scales as *N*.

Industry Trends

ATM is widely used by commercial service providers for a variety of purposes, such as transporting Internet traffic in service provider cores and supporting digital subscriber line service (which provides high-speed connections to the Internet using home telephone lines). Revenues from providing ATM services in the U.S. market have grown steadily since the late 1990s, increasing more than tenfold (in nominal dollars) since 1997. Annual sales of ATM equipment, however, rose through 2000 and then fell by more than 50 percent through 2003. A key reason for that drop was the decline in the fortunes of the U.S. telecommunications industry that began in 2000. The consensus market forecast appears to be that revenues from ATM services will continue to grow fairly strongly and that sales of ATM equipment will halt their decline but then grow only modestly. Many observers cite the emergence of competitors to ATM, particularly MPLS, as the reason for that modest growth.[25]

Internet protocol traffic traversing commercial core networks has grown more than sixfold since 2000. (However, revenues from providing Internet services have declined because of competition.) Internet protocol traffic, instead of voice, now accounts for the majority of traffic. Because MPLS is well suited to provide transport of converged services using the Internet protocol, many observers forecast that sales of MPLS-capable equipment will increase, perhaps strongly. An indication of the future course of the market may be that the separate industry associations formed to develop standards for and promote the use of ATM and MPLS—the ATM Forum and the MPLS and Frame Relay Alliance—announced in July 2004 that they plan to merge.

---

25. Telecommunications Industry Association, *2004 Telecommunications Market Review and Forecast* (Arlington, Va.: TIA, 2004).

Capacity of the GIG-BE

One measure of a network's capacity is the total "raw" bandwidth (not accounting for overhead) provided in its core to transport traffic of all kinds. At its full operational capability in September 2005, the GIG-BE core—which will consist of 66 service-delivery nodes in the continental United States—will have an average raw capacity per service-delivery node of 120 Gbps. That size is about 21 times larger than the average capacity per core node of the DATMS (5.6 Gbps for the 20 core nodes in the continental United States).[26] In addition, the equipment used in the GIG-BE core for optical transport of data can be upgraded to provide an average capacity per core node of 945 Gbps.

Another measure of capacity is the so-called tributary bandwidth provided to users to feed traffic into a network. The DATMS network in the continental United States contains 230 ATM tributary switches that forward data to its 20 core nodes. The total capacity of the links connecting those switches is 104 Gbps, implying an average tributary capacity of 454 Mbps. The GIG-BE, in contrast, is not composed of distinct tributary and core segments. However, at the 46 sites in the continental United States where a GIG-BE service-delivery node coincides with a tributary site of the DATMS, the GIG-BE will increase the available capacity by more than a factor of 260—that is, a 120 Gbps average capacity at all GIG-BE core nodes compared with the 454 Mbps capacity at DATMS tributary sites.

The total bandwidth provided to transport traffic of all kinds is not the same as the useful bandwidth available to users to transport their data. A number of factors make useful bandwidth smaller than total bandwidth, including the overhead associated with network operations. That overhead includes Internet protocol headers, MPLS labels (also called shim headers), ATM cell headers, and other things associated with ensuring the network's security. An additional source of inefficiency in transporting Internet traffic over ATM is the need to segment the Internet protocol's variable-length packets so they fit into the 48 bytes available to transport data in an ATM cell.

For the transport of Internet traffic that is not encrypted, MPLS is 18 percent more efficient than ATM (see Table 2). For encrypted traffic, however, that advantage is reduced, and under certain conditions, ATM can be the more efficient alterna-

---

26. That comparison does not account for differences in the efficiency with which the GIG-BE and the DATMS transport traffic (those differences are considered later in this review). The GIG-BE has 19 nodes that are located at corresponding DATMS core nodes in the continental United States. The average capacity of the GIG-BE at those 19 nodes will be 154 Gbps, 27 times greater than the corresponding capacity of the DATMS (5.6 Gbps). Both the DATMS and GIG-BE have nodes located outside the continental United States. The total raw capacity of the GIG-BE at its full operational capability will be 7,878 Gbps, about 70 times greater than the total raw capacity of the DATMS (113 Gbps).

Table 2.

Ratio of MPLS Efficiency to ATM Efficiency for Transporting Internet Protocol Version 4 Traffic

	Encrypted Traffic		
		HAIPE—Tunnel Mode	
Unencrypted Traffic	HAIPE—Transport Mode	Without Traffic-Flow Security Padding	With Traffic-Flow Security Padding
1.18	1.10	1.06	0.74

Source: Congressional Budget Office using a sample of Secret Internet Protocol Network traffic provided by the Defense Information Systems Agency.

Notes: MPLS = multiprotocol label switching; ATM = asynchronous transfer mode; HAIPE = high-assurance Internet protocol encryption.

These estimates assume no use of header compression or bandwidth-saving techniques. Overhead is assumed to be any transmitted information (other than user data) that is added at layer 2 or higher of the network. The estimates assume Internet protocol encryption block length of eight bytes, IPv4 header of 20 bytes, other overhead composing 20 bytes, an MPLS label stack comprising eight bytes (two labels deep), and use of ATM adaption layer 5. The use of bandwidth-saving techniques would increase the ratios for transporting encrypted traffic shown here by about 4 percent.

tive. The reason is that the methods used to encrypt ATM cells add less overhead than the current version of high-assurance Internet protocol encryption (HAIPE) that will be used in conjunction with MPLS by the GIG-BE. Although more-efficient versions of HAIPE are being developed, the current version adds about 40 bytes of overhead in what is called transport mode, about 60 bytes (for Internet protocol version 4 traffic) in what is called tunnel mode, and up to many hundreds of bytes (depending on packet size) in tunnel mode with padding to ensure traffic-flow security.[27] The additional overhead associated with each of those three modes of operation provides additional security. For example, in transport mode, an Internet header containing information that could be used to infer the structure of the classified networks feeding traffic into the GIG-BE is sent unencrypted. In tunnel mode, that header is encrypted, and a new header, potentially revealing information only about the structure of the GIG-BE's core, is added and sent

27. National Security Agency, *Interoperability Specification for High Assurance Internet Protocol Encryptor Devices*, version 1.3.4 (October 14 , 2003). Future versions of the specification are being developed with bandwidth-saving options that could reduce overhead to 18 bytes in transport mode and to 38 bytes in tunnel mode for Internet protocol version 4 traffic. As of October 2004, version 3.0.0 of the HAIPE specification incorporating those bandwidth-saving modes was expected to be released in the spring of 2005.

unencrypted.[28] Traffic-flow security padding obscures the size of the data packets that users are sending and is used only in situations in which it is thought to be possible that untrusted outsiders could monitor a network's traffic. That possibility is considered unlikely for the GIG-BE core network in the continental United States, so the use of traffic-flow security padding for GIG-BE traffic is not required.[29]

Another factor that affects the GIG-BE's capacity to transport classified data is the speed of the high-assurance Internet protocol encryptor devices now available. Initially, the GIG-BE will use HAIPE devices that operate at 100 Mbps, with an effective throughput of about 80 Mbps. Certification of one of those devices by the National Security Agency (NSA) was anticipated in September 2003 but was delayed until February 2004. Two HAIPE devices for use in the GIG-BE that operate at 1 Gbps, with an effective throughput of 800 Mbps to 900 Mbps, are under development, and one is being testing by the Defense Information Systems Agency. The NSA is expected to certify those devices by late February 2005. In addition, a HAIPE device operating at 10 Gbps is being developed; as of October 2004, its NSA certification was anticipated by November 2005. An ATM cell encryption device operating at 10 Gbps was certified by the NSA for use in December 2004.[30]

Accounting for overhead and for the throughput of the available HAIPE devices, the Congressional Budget Office estimates that the additional capacity to transport encrypted traffic provided initially by the GIG-BE using 100-Mbps HAIPE devices will be about double the capacity of the DATMS (see Table 3). When 1-Gbps HAIPE devices are employed, the additional encrypted capacity provided by the GIG-BE will be about 10 times the capacity of the DATMS. When the GIG-BE is at its full operational capability (projected for September 2005), the increase in total capacity—both encrypted and unencrypted—will be about 80 times the capacity of the DATMS, CBO estimates. Moreover, the possibility

---

28. Protecting information that reveals the structure of the edge networks that generate classified traffic is one example of a feature associated with traffic-flow security. However, the use of tunnel mode can be necessary in many instances for reasons other than traffic-flow security. For example, if a site connects to the GIG-BE using two access paths, the use of tunnel mode is necessary to ensure that classified traffic is forwarded to the appropriate HAIPE device for decryption.

29. For transport outside the continental United States, the GIG-BE will use bulk encryption at the link layer.

30. Making full use of the capacity of the 10-Gbps ATM cell encryptor would, however, require its use in an ATM network with switches operating at speeds of at least 10 Gbps and populated with 10-Gbps encryptors. The ATM switches currently used in the DATMS core operate at 622 Mbps, and the majority of ATM cell encryptors now deployed operate at speeds of 654 Mbps. Other government networks than the DATMS exist that have recently installed high-speed ATM switches and that could make full use of 10-Gbps ATM encryptors.

Table 3.

Capacity Provided by the GIG-BE Relative to That of the DATMS

(Multiples of DATMS capacity)

	Using 100-Mbps HAIPE Devices[a]	Using 1-Gbps HAIPE Devices[a]	Maximum Possible with Upgrades to Current Equipment[b]
Total Available Bandwidth	80	80	650
Encrypted	2	10	70
Unencrypted	115	110	900

Source: Congressional Budget Office based on data from the Defense Information Systems Agency.

a. Encrypted capacity is estimated assuming that four HAIPE devices would be used at each of the GIG-BE's 66 core service-delivery nodes. If more HAIPE devices were employed at each node, the encrypted capacity provided by the GIG-BE would be larger than shown here.

b. Assuming that the fractions of encrypted and unencrypted traffic associated with use of the 1-Gbps HAIPE devices in the GIG-BE as configured at its full operational capability would be the same if the network was upgraded. That assumption implies a substantial increase in encrypted capacity that might require development and use of HAIPE devices with a throughput substantially greater than those now existing if the use of multiple sets of existing HAIPE devices proved infeasible.

exists to increase that capacity by about another factor of eight if the GIG-BE's equipment is upgraded to the maximum extent possible.

Some projections by the Defense Information Systems Agency assume that the need for capacity will grow at annual rates of 50 percent in the future. Given the total raw capacity now provided by the DATMS, at that rate of growth, the core capacity provided by the GIG-BE at its full operational capability would not be outstripped until about 2015 (see Figure 1). Projected needs would not exceed the full potential capacity of the GIG-BE with possible upgrades until about 2020.

Security of the GIG-BE

The approach used in the GIG-BE to encrypt classified traffic and separate it from unclassified traffic differs from the approach used in DoD's existing networks. The GIG-BE employs high-assurance Internet protocol encryption performed at layer 3 and virtual private networks created using MPLS to separate classified and unclassified traffic, allowing both types of traffic to be transported over a common layer-1 infrastructure. DoD's existing networks, such as the DATMS, have employed virtual private networks created at layer 2 using ATM, as well as ATM cell encryption, to allow classified and unclassified traffic to be transported over a common layer-1 infrastructure. The unclassified networks supported by both the

Figure 1.

Total GIG-BE Capacity Compared with Projected Needs

(Raw capacity in gigabits per second)

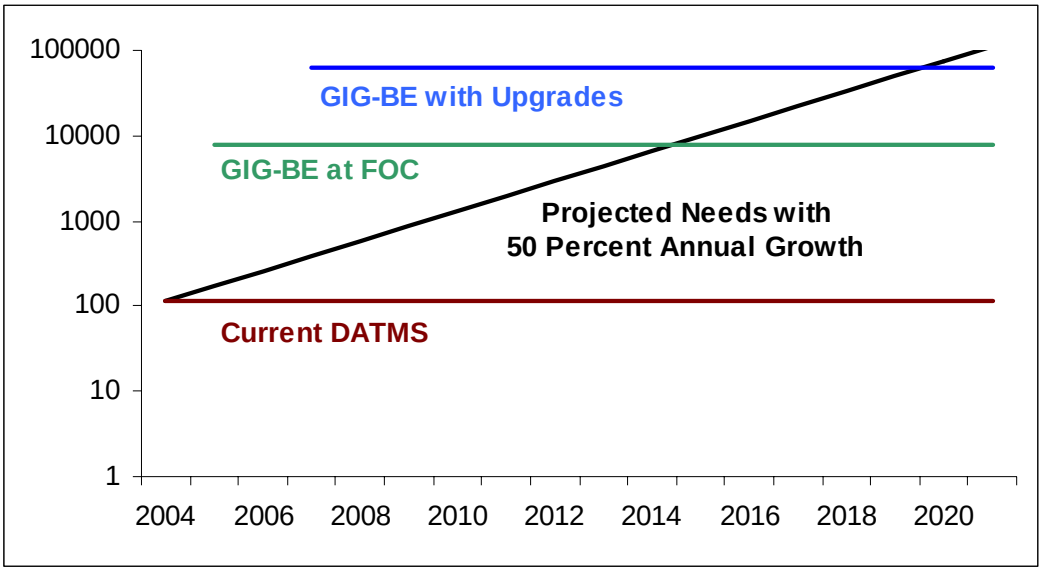

Source: Congressional Budget Office.

Note: GIG-BE = Global Information Grid Bandwidth Expansion; FOC = full operational capability; DATMS = DISN (Defense Information System Network) Asynchronous Transfer Mode System.

GIG-BE and the DATMS are connected to the Internet (indirectly through the Non-Classified Internet Protocol Router Network and behind firewalls).[31]

Virtual Private Networks

As noted above, a virtual private network (VPN) comprises a set of users at different sites, often geographically dispersed, that communicate with each other using a common core network such as the GIG-BE. The GIG-BE will use MPLS label switched paths—as opposed to ATM virtual circuits—to establish separate VPNs for users sharing classified or unclassified traffic.[32] When the virtual private networks are properly configured, traffic originating on one VPN cannot be seen

31. Firewalls are hardware devices or software programs designed to protect computer systems from unauthorized access. Network firewalls are generally used to protect an internal computer network from outside access (although they can also be used to restrict internal users' access to the outside). In addition, they can hide the network's structure from outsiders. Firewalls can receive traffic at the interface between networks and block it as necessary.

32. Internet Engineering Task Force, *Request for Comments: 2547—BGP/MPLS VPNs* (March 1999), available at www.ietf.org/rfc.html. (BGP stands for border gateway protocol, a method used by routers in the GIG-BE to exchange routing information with routers in the networks composing the VPNs outside the GIG-BE core.)

by users on another VPN—a situation sometimes referred to as logical separation. Thus, VPNs keep traffic "virtually private" even though it is being transported over a common core network, or physical infrastructure.

Both ATM and MPLS can be used to support such virtual private networks, but with a difference that some observers see as significant: MPLS uses layer-3 Internet protocols for control, whereas ATM uses unique control protocols and generally operates at layer 2. (ATM does use some protocols operating at layer 3, however, to establish and deestablish connections.) Some observers have argued that MPLS's use of layer-3 protocols makes virtual private networks created with MPLS vulnerable to denial-of-service attacks launched from other VPNs or from the Internet. However, according to the Internet Engineering Task Force's guidance document for MPLS virtual private networks, "a virtual private network should, even without the use of cryptographic security measures, provide a level of security equivalent to that obtainable when a layer 2 backbone . . . is used. That is, in the absence of misconfiguration or deliberate interconnection of different VPNs, it should not be possible for systems in one VPN to gain access to systems in another VPN."[33]

Virtual private networks running over any core, including the GIG-BE, are vulnerable to attacks launched from within that core if it is not secure. If attacks from within the core are a concern, a primary way to increase security is to employ encryption (regardless of whether the core uses ATM, MPLS, or some other technology).

Encryption

Cryptographically enforced VPNs—which the GIG-BE will create using high-assurance Internet protocol encryption—provide protection against traffic monitoring and misconfiguration of switches and routers that the so-called logical VPNs described above (which the GIG-BE will also use) do not provide.[34] In addition to traffic encrypted with HAIPE, the GIG-BE will transport other encrypted traffic (including traffic using ATM cell encryption) generated by users of DoD's legacy networks.

HAIPE devices incorporate a number of modes of operation, which provide different kinds of protection. (Some of those modes, such as tunnel and transport,

---

33. Internet Engineering Task Force, *Request for Comments: 2547—BGP/MPLS VPNs*. That document was developed in the routing area of the IETF, however, which is not composed of the organization's experts on security. Thus, some observers would argue that its claims about security should not be viewed as definitive (or in all cases correct).

34. National Security Agency, *Interoperability Specification for High Assurance Internet Protocol Encryptor Devices*, version 1.3.4.

were described above in the section about the capacity of the GIG-BE.) In general, the HAIPE specification requires that the default modes of operation for HAIPE devices be those offering the most protection—and imposing the greatest overhead. Users can choose to operate their networks in other than those HAIPE default modes, subject to guidance promulgated by the National Security Agency.[35] In general, links that are located within the United States in networks managed by government agencies are considered to be at low risk of compromise and do not require the use of a number of the most secure modes of HAIPE operation. The GIG-BE core located in the continental United States is such a network.

In May 2003, when DoD announced that the GIG-BE would be MPLS-based and would use high-assurance Internet protocol encryption, the fastest HAIPE devices operated at speeds about a factor of 25 slower than the fastest ATM encryption devices. Although a HAIPE device operating at 1 Gbps is expected to be certified by the NSA in February 2005, it is slower than the most recent ATM encryptor, a 10-Gbps FASTLANE model certified in December 2004. NSA currently expects to certify a 10-Gbps HAIPE device in November 2005. However, to take advantage of higher-speed encryptors, a network must have switches or routers that operate at least as fast as the encryptors, and many of the network's users who transmit encrypted traffic must be using those higher-speed encryptors . The majority of the ATM cell encryptors bought within the past several years operate at 654 Mbps, consistent with the operating speeds of the ATM switches used in the DATMS. Classified networks other than the DATMS, however, have reportedly been upgraded in the past year with high-speed ATM-capable combination switch-routers that could be used in conjunction with the 10-Gbps model of the FASTLANE, provided that the users on those networks bought and operated that faster encryptor. DoD does not expect to buy ATM encryptors with speeds exceeding the 10-Gbps FASTLANE now available.

Debate continues among members of the national security community about the benefits and drawbacks of using link encryption at layer 1, ATM encryption at layer 2, or high-assurance Internet protocol encryption at layer 3. An assessment by the staff of the NSA came to the following conclusions:[36]

- When link encryption is possible to implement, it provides the best security because it offers no visibility into the underlying structure of the traffic

35. National Security Agency, *Internet Protocol Configuration Guidance for High Assurance IP Encryptor Networks*, version 2.0 (October 22, 2004).

36. Chris Kubic, Technical Director for Architectures and Technology, Network and Space Information Assurance Office, National Security Agency, "ATM/IP Encryption Comparison" (briefing).

being transported. However, it does not allow unused bandwidth to be used for transporting traffic at other levels of security.[37]

- ATM encryption is less complex than high-assurance Internet protocol encryption because ATM encryptors operate on fixed-length cells instead of variable-length Internet protocol packets. ATM encryption also imposes minimal overhead; but its management is complex, and its use constrains the ways in which networks can be configured and operated.

- The use of high-assurance Internet protocol encryption provides "converged security" in that it operates at layer 3 and does not use ATM's separate and unique infrastructure. However, it is much less efficient than ATM for transporting small packets (which compose voice traffic), and its traffic-flow security features have "big bandwidth implications."[38]

- Both ATM encryption and high-assurance Internet protocol encryption are problematic because of their use of unencrypted headers.

- NSA does not have a preference for either method of encryption.

- The challenges associated with providing security for traffic transported on a converged Internet protocol backbone overshadow the issue of whether to use HAIPE or ATM encryption.[39]

In addition to those points, some observers would note that ATM cell encryption performed at layer 2 is sensitive to disrupted operation of a core network in ways that high-assurance Internet protocol encryption performed at layer 3 is not.[40]

Issues have also been raised about the strategy that DoD is developing to provide information assurance for the overall GIG (of which the GIG-BE is a compo-

---

37. Link encryption can be used with ATM, MPLS, or any of the other common methods for forwarding traffic over a backbone network.

38. Some network engineers would note that under certain circumstances, ensuring traffic-flow security using ATM cell encryption has similar implications for network overhead and available bandwidth as using traffic-flow security padding with HAIPE devices.

39. This review does not directly address those challenges.

40. ATM cell-encryption devices associate encryption keys with individual virtual circuits. If the operation of an ATM switch is disrupted, all of the virtual circuits starting and ending at that switch must be reestablished, which means that the associations between circuits and encryption keys must also be reestablished. High-assurance Internet protocol encryption associates encryption keys with connection endpoints. Thus, when router operations in the core network are disrupted, key associations do not have to be reestablished.

nent).[41] In December 2004, a review group formed by the NSA and composed of industry experts concluded the following about DoD's approach to providing information assurance for the overall GIG:

- A detailed architecture defining the GIG in the near term is nonexistent;
- Plans that do exist for implementing the GIG are problematic, requiring management of vast quantities of information never attempted before and depending on technologies that do not exist and may not be feasible;
- The GIG's survivability and robustness have not been addressed; and
- "Risk is unbounded within the GIG vision."[42]

None of those conclusions—which focus on DoD's plans for the GIG as a whole—would have been different if ATM had been chosen instead of MPLS as the technology for forwarding traffic in the GIG-BE.

The GIG-BE's Quality of Service

Quality of service is the performance observed by a user on a network. Key measures of performance include the time delay between when data are sent and received, the degree to which that delay varies, and the extent to which data are lost. Transmission of different types of data requires different kinds of quality of service. For example, transmission of voice in real time can tolerate some loss but requires short delay and little variation in delay. Transmission of data such as e-mail messages, by contrast, can tolerate substantial delay and variation in delay.

ATM Quality of Service

Asynchronous transfer mode assigns numerical values to parameters that determine required performance. ATM virtual circuits are established and traffic is assigned to them on the basis of those performance requirements. The ATM

41. "Information assurance" refers to measures taken to protect and defend information and information systems by ensuring their availability, integrity, authentication, confidentiality, and nonrepudiation. Availability is the extent to which a system is operable and capable of performing its intended functions. Integrity is the extent to which the system and associated data are reliable and logically correct. Authentication ensures that users and their transmissions are valid and authorized. Confidentiality ensures that information is not revealed to people not authorized to obtain it. Nonrepudiation guarantees that a message or data can be proved to have originated from a specific person.

42. William Neugent, "Senior Industry Review Group Recommendations and Observations" (briefing prepared for the National Security Agency, December 9, 2004).

Forum defines six categories of service that can be associated with an ATM connection:[43]

- *Constant bit rate*, which is used for traffic that needs a specified amount of bandwidth to be available as long as the connection is maintained. Voice and video transmitted in real time are examples of traffic that might need such service.

- *Real-time variable bit rate*, for traffic that is sensitive to delay and variation in delay, such as voice and video, but that is transmitted at varying rates and therefore does not need a guaranteed amount of bandwidth.

- *Non-real-time variable bit rate*, for traffic, such as the transfer of data files, that can require varying amounts of bandwidth and is not sensitive to delay.

- *Unspecified bit rate*, for non-real-time traffic, such as e-mail messages, that is not sensitive to delay and for which delivery does not need to be guaranteed. Such traffic is given "best-effort" service in a manner similar to that provided by many Internet protocols.

- *Available bit rate*, for transmitting non-real-time traffic at rates that can be continually adjusted on the basis of demands on the network.

- *Guaranteed frame rate*, for non-real-time traffic that needs a minimum amount of bandwidth.

Within each category of service, two sets of numerical parameters are specified: one set, called traffic parameters, characterizes the traffic that users expect to transmit over the connection; the second set defines the quality of service to be provided for transmitting traffic that is consistent with the specified traffic parameters. The six traffic parameters include the peak and minimum rates at which users will transmit ATM cells over a connection. The six quality-of-service parameters include the fraction of cells transmitted that can be lost and the maximum acceptable delay in transmitting a cell.[44] The user and provider enter into service-level agreements in which values can be specified for those parameters for each type of traffic that the user expects to transmit. Those agreements allow quality of service tailored to specific types of traffic to be strictly guaranteed.

---

43. ATM Forum Technical Committee, *Traffic Management Specification Version 4.1*, AF-TM-0121.000 (March 1999).

44. Only three of the six quality-of-service parameters can be negotiated (that is, specified by agreement between a user and a network service provider).

The specifications for the DATMS state that it should support the use of all of the ATM quality-of-service features described above.[45] However, CBO does not have information indicating how, as a practical matter, those requirements have been implemented in the day-to-day operations of the DATMS.

MPLS Quality of Service

A key feature of multiprotocol label switching that affects the quality of service it can be used to provide and that distinguishes it from the Internet protocol is MPLS's ability to support traffic engineering (optimizing the use of network resources by transmitting traffic on particular paths). Traffic engineering involves computing a path on which traffic can flow from source to destination that is subject to a set of constraints. Such engineering is not possible with the Internet protocol because it does not predetermine end-to-end paths for traffic to traverse; rather, the path that traffic travels is determined incrementally at each intermediate node in the network solely on the basis of the packet's intended destination.

Currently, networks (such as the GIG-BE) that use MPLS can provide users with many of the quality-of-service features of ATM. The reasons for that are the use of the Internet's architecture for differentiated services in conjunction with an extension to the Internet's resource reservation protocol that MPLS uses to establish its label switched paths.[46] Differentiated services (diffserv) assigns relative priorities that differentiate how packets are processed as they transit a network, using up to six bits in the Internet protocol header.[47] The indicator assigned to a packet is called a diffserv code point; its value can depend on both the type and priority of the traffic. A working group of the Military Communications Electronics Board (MCEB) has defined seven types of traffic and up to five levels of precedence for each traffic type, requiring 25 diffserv code points. Traffic types include interactive (real-time) voice messages, interactive video, and batch transfer of long blocks of data. Precedence levels comprise flash override, flash, immediate, priority, and routine.

---

45. Department of Defense, *Defense Information System Network Asynchronous Transfer Mode System Specification Version 1.2c* (April 17, 1998).

46. Internet Engineering Task Force, *Request for Comments: 2475—An Architecture for Differentiated Services* (December 1998), *Request for Comments: 3209—Extensions to RSVP for LSP Tunnels* (December 2001), and *Request for Comments: 3270—Multi-Protocol Label Switching Support of Differentiated Services* (May 2002), available at www.ietf.org/rfc.html.

47. David E. McDysan and Dave Paw, *ATM & MPLS Theory and Application: Foundations of Multi-Service Networking* (Emeryville, Calif.: McGraw-Hill/Osborne, 2002). However, methods have also been defined to assign and measure numerical performance parameters—rather than relative priorities—for Internet protocol and MPLS forwarding; see International Telecommunications Union, *Internet Protocol Data Communication Service—IP Packet Transfer and Availability Performance Parameters*, Recommendation Y.1540 (November 13, 2002), and *Performance and Availability Parameters for MPLS Networks*, Recommendation Y.1561 (February 13, 2004), available at www.itu.int/ITU-R/publications/rec/index.asp.

MPLS can use the diffserv code points in two ways:

- To determine up to eight ways in which a packet assigned to any label switched path should be processed by routers as the packet transits the network (three bits are available in the MPLS label for that purpose); and

- To determine the label switched paths to which traffic is assigned, as well as how the traffic should be processed by the network's routers.

Initially, the GIG-BE will use a simplified version of the first approach. Its label edge routers will assign packets carrying one of the MCEB's 25 diffserv code points to one of four classes of service for forwarding through the GIG-BE's core. According to DoD, that approach is being used to simplify network management and control, consistent with current practices in large networks established by commercial Internet service providers.[48] That approach does not provide all of the quality-of-service capabilities of ATM.

MPLS, using diffserv and extensions of the Internet's resource reservation protocol (and other routing protocols), can establish label switched paths based on the bandwidth needed and available to transmit traffic. But the bandwidth constraints are applied at an aggregate level across all of the classes of service into which the traffic that is flowing has been divided; bandwidth guarantees are not made for specific types of traffic or classes of service.[49] An approach has been defined, however, for developing label switched paths and assigning traffic to them according to the bandwidth needs of specific types of traffic, such as voice.[50] When that approach is implemented in MPLS networks such as the GIG-BE, it will allow quality of service to be strictly guaranteed for specific types of traffic, in a manner analogous to that of ATM. The latest versions of the operating system used in the routers purchased for the GIG-BE have that capability (called diffserv-aware traffic engineering), although it has not yet been implemented in the network.

In general, the larger the number of service-associated parameters that can be assigned, processed, and measured, the more complex network operations will be, but the more stringently quality of service can be controlled. Some people regard ATM networks as complex to operate, but they do not dispute that ATM can

---

48. David Mihelcic, "GIG Bandwidth Expansion Update" (briefing to Congressional staff, February 18, 2004).

49. Juniper Networks, *MPLS DiffServ-Aware Traffic Engineering*, white paper (Sunnyvale. Calif.: Juniper Networks, 2004).

50. Internet Engineering Task Force, *Request for Comments: 3564—Requirements for Support of Differentiated Services-Aware MPLS Traffic Engineering* (July 2003).

ensure good quality of service. However, in networks where bandwidth is plentiful—as is likely to be the case for GIG-BE operations initially and probably for a number of years thereafter—good quality of service can be provided while using simpler methods because there is not the competition for limited bandwidth that can make stringent control necessary.

The Potential for the GIG-BE to Benefit the Commercial Sector

While testifying to the Congress in early 2003 (before the initial requests for proposals for the GIG-BE were issued), the Assistant Secretary of Defense for Networks and Information Integration discussed the potential for the expansion program to benefit the commercial sector.[51] That issue arose in part because of the large increase in capacity that the GIG-BE would provide relative to DoD's legacy networks. In particular, if the network was upgraded to the full capabilities of the equipment now being purchased, the GIG-BE would have more than 650 times the capacity of DATMS. Some observers considered such an increase large enough that building the GIG-BE might significantly benefit the depressed U.S. telecommunications industry.

DoD's total investment in the GIG-BE, however, has not been large relative to annual U.S. commercial telecommunications investment for the past several years (see Table 4). Moreover, the amount of dark fiber-optic cable that the GIG-BE will use is a small fraction of the total available. Thus, it seems unlikely that the GIG-BE will have a significant effect in the short run on the commercial telecommunications industry in the United States.

51. Testimony of John P. Stenbit before the Subcommittee on Terrorism, Unconventional Threats, and Capabilities of the House Armed Services Committee, April 3, 2003.

Table 4.

The GIG-BE Compared with U.S. Telecommunications Networks

	GIG-BE	U.S. Commercial Telecommunications
Spending for Networks	~$900 million total	~$300 billion annually[a]
Spending for Routers	~$70 million total	~$2 billion annually[a]
Dark Fiber	<100,000 miles lit[b]	15 million miles unlit[c]

Source: Congressional Budget Office.

a. Gartner Dataquest, December 2003.

b. The Congressional Budget Office's upper-bound estimate of the amount of unused dark fiber that is intended to be "lit up" by the Global Information Grid Bandwidth Expansion.

c. Fred Donovan, *Fiber-Optic News*, 2001.

Appendix: Network Concepts

This appendix discusses basic network concepts, including the distinctions between the core and edge of a network, the different layers that compose a network, the distinctions between a network's data and control planes, and the different types of data that are transmitted over networks.

Network Core and Network Edge

The core of a network is the set of nodes connected by the highest-capacity communications links along which multiplexed data are transmitted and directed using the highest-capacity switches and routers (see Figure A-1). Data moving into or out of the network core are managed by edge routers or switches (see Figure A-2). Local area networks (LANs) physically located at sites worldwide connect to the core network through its edge routers or switches (see Figure A-3).

Network Layers

The Open Systems Interconnection Reference (OSIR) Model, which was published in 1983 by the International Organization for Standardization, defines the functions and protocols that enable computer systems to connect to one another. The model consists of seven layers, each with specific functions.

1. *Physical layer*—handles the interface to the physical medium used to transmit data, such as fiber-optic cable, copper wires, or radio waves. It is where information is transmitted as a stream of bits.

Figure A-1.

Core Network with Routers or Switches at the Nodes

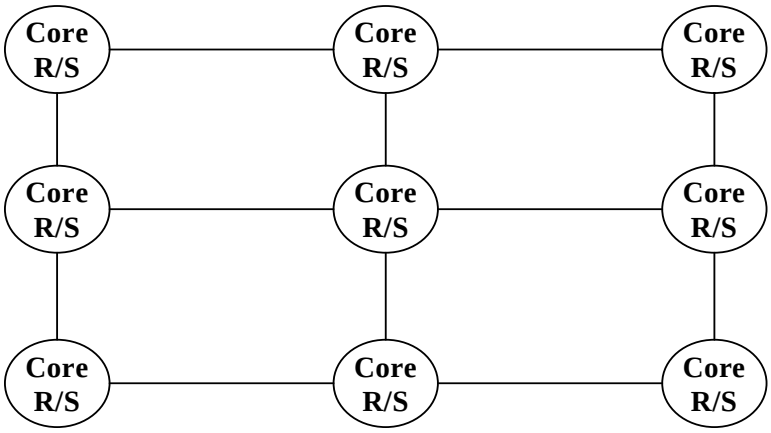

Source: Congressional Budget Office.

Note: R/S = router or switch.

Figure A-2.

Schematic of a Routing or Switching Network

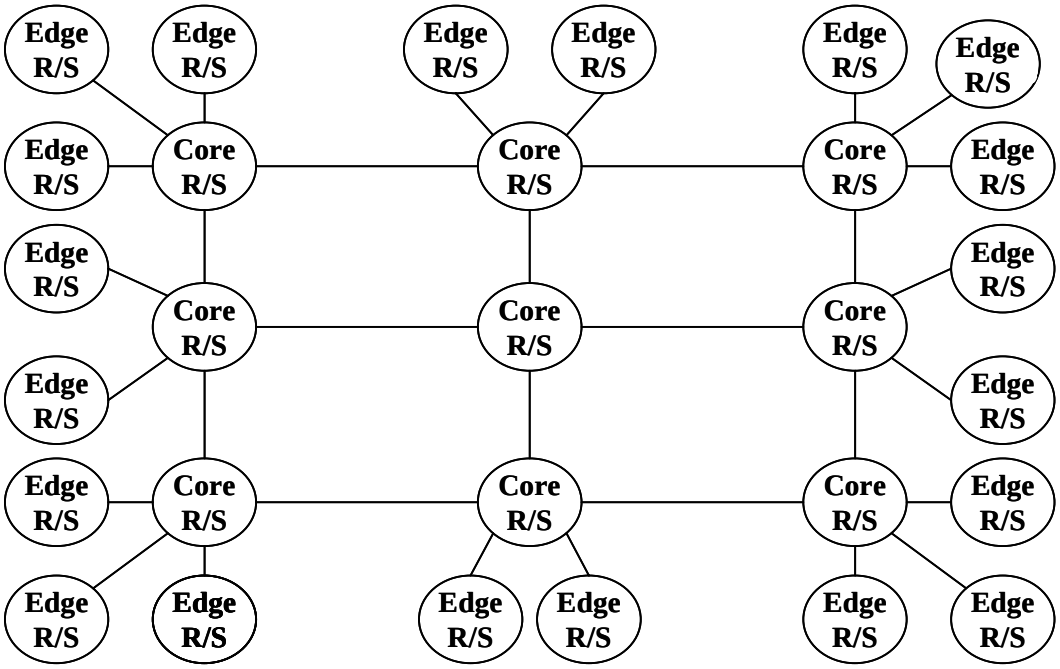

Source: Congressional Budget Office.

Note: R/S = router or switch.

2. *Data link layer*—can provide error-free transmission for the network layer and controls the flow and delivery of data (for example, packet forwarding or cell transmission) over a network's connections.

3. *Network layer*—handles routing and congestion; also establishes and tears down connections between devices communicating over a network.

4. *Transport layer*—interconnects entities operating at layers 1 through 3. It can control the flow of, and detect and correct errors in, data transferred between communicating devices.

5. *Session layer*—provides the user's interface to the network. For example, it provides connections between a workstation connected to a LAN and a central processor.

Figure A-3.

Schematic of a Full Network with Core and Edge Routers or Switches, LANs, and Users

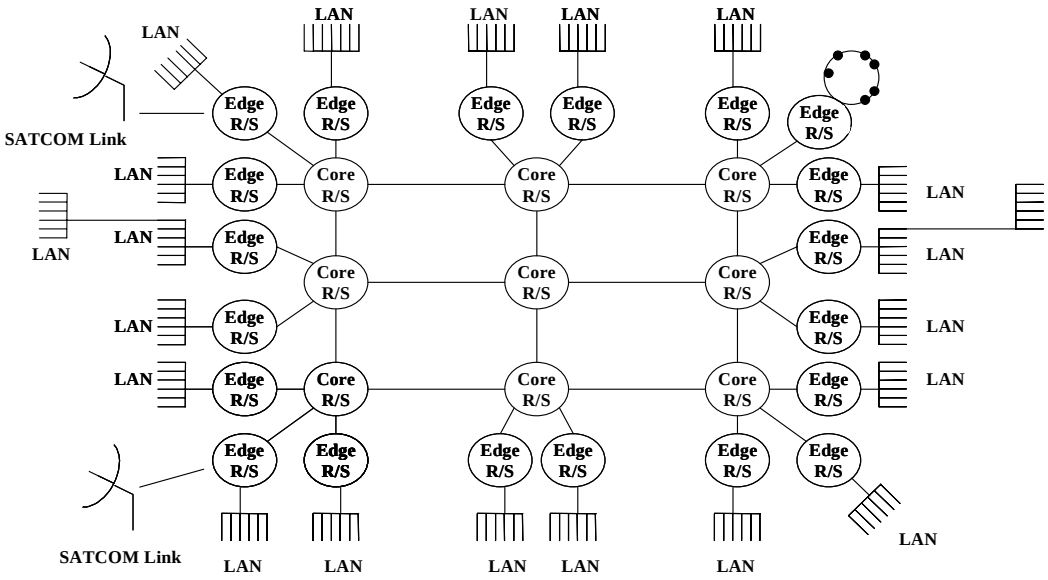

Source: Congressional Budget Office.

Notes: SATCOM = satellite communications; LAN = local area network; R/S = router or switch.

In this example, there are nine core nodes (either routers or switches), 20 edge nodes (either routers or switches), 22 LANs (with two being ethernet LANs connected over bridges and one being a token ring LAN), two SATCOM links, and 133 users explicitly connected via the LANs.

6. *Presentation layer*—determines how data is displayed to a user through video and text display formats.

7. *Application layer*—manages the software and hardware generating the data that is to be transported over the network. This is the layer with which end users interact directly.

The OSIR model is both general and idealized, and not all of its layers and features apply to or are used in communications networks operating today. For example, the Internet uses only the application, transport, and network layers (it embeds services of the presentation layer in the application layer).

Data Plane Versus Control Plane

The transmission of a single e-mail message over a network may be accompanied by the generation and transmission of many other messages. Those other messages include acknowledgments that the e-mail has arrived at a given node and notifications that errors have been detected as the e-mail is sent between nodes. Such messages, which contain the information that the network needs to operate itself, are transmitted on what is termed the "control plane." The sender and receiver of the e-mail do not see those messages; they see only messages transmitted on the network's data (or user) plane. Control-plane functions include establishment, monitoring, and deestablishment of virtual circuits in asynchronous transfer mode and label switched paths in multiprotocol label switching.

Different Types of Data

Many different types of data are transmitted among a network's users; those data can usually be distinguished by how quickly or accurately they must be received. Some data—such as voice transmissions (that is, telephony)—are time-sensitive and must be transmitted quickly, though not necessarily with perfect accuracy. For other data—such as e-mail messages—accurate transmission may be more important than quick receipt.

In general, data are broken into many pieces, or packets, for transmission over a network. Those packets must be reassembled in order for the full set of data being sent to be received correctly. Consequently, the speed and accuracy with which the packets are transmitted and reassembled determines the speed and accuracy with which the overall set of data is received.

Two standard measures are used to define time sensitivity: latency and jitter. Latency is the time delay between when a set of data is sent and when it is received. Jitter is the variation in latency among the packets into which the data are decomposed for transmission over the network. The amounts of latency and jitter that are acceptable depend on the type of data being transmitted. For example, people expect a typical voice data stream to be essentially continuous and received nearly immediately by the hearer (that is, within one to two tenths of a second of the words' being spoken). Long latency in voice communication is generally unacceptable because it can induce both speakers to talk at the same time (the sender because he has not heard an expected response from the recipient, who also thinks the sender has stopped talking). For voice communication, latency on the order of several tenths of a second is too long for the quality of transmission to be considered good. For e-mail messages, however, latencies on the order of seconds to minutes are usually acceptable.

In terms of jitter, the greater it is, the more a continuous voice conversation seems to be broken up into unnatural chunks. For example, jitter can occur in mobile cell phone use near the edge of service-area boundaries, when the relay stations being used to receive or transmit the cell phone's signal may be changing. That change causes variance in the time required to transmit the sequential packets into which the phone's transmissions are decomposed. In general, voice communication has the highest sensitivity to both latency and jitter.

Although telephony and teleconferencing are sensitive to latency and jitter, they are relatively insensitive to noise. Because of the processing capability of the human brain, people can hear and see properly under widely varying conditions of background noise and lighting. Thus, variations of a few percent in the amplitude of received audio and video signals are often not even recognized.[1]

Unlike audio and video communications, numerical data files—for example, payroll files—can require a high degree of accuracy for a successful transmission. If a network routinely transmitted a file containing biweekly pay authorizations and introduced errors on the order of 2 percent, it would most likely be considered unacceptable. By contrast, today's data-network users generally do not require e-mails, faxes, or file transfers to occur instantaneously. Latencies on the order of seconds or minutes are usually acceptable because such messages are often simply stored for later review.

1. When voice (or other data) are transmitted in packets, noise as described here is usually detected and rejected at layer 2. But when unreliable delivery of packets occurs, packets can be dropped. In that case, the techniques generally used to transmit voice and video (called compression) require a receiver to interpolate the content of lost packets using the data in packets received successfully. That approximation results in an effect equivalent to noise at the receiver.